arXiv:1606.07425v2 [cs.DS] 27 Jun 2016

Generalized Preconditioning and Network Flow Problems

Jonah Sherman*
University of California, Berkeley

June 28, 2016(preliminary draft)

Abstract

We consider approximation algorithms for the problem of finding x of minimal norm $\|x\|$ satisfying a linear system $\mathbf{A}x = \mathbf{b}$, where the norm $\|\cdot\|$ is arbitrary and generally non-Euclidean. We show a simple general technique for composing solvers, converting iterative solvers with residual error $\|\mathbf{A}x-\mathbf{b}\| \leq t^{-\Omega(1)}$ into solvers with residual error $\exp(-\Omega(t))$, at the cost of an increase in $\|x\|$, by recursively invoking the solver on the residual problem $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{A}x$. Convergence of the composed solvers depends strongly on a generalization of the classical condition number to general norms, reducing the task of designing algorithms for many such problems to that of designing a *generalized preconditioner* for $\mathbf{A}$. The new ideas significantly generalize those introduced by the author's earlier work on maximum flow, making them more widely applicable.

As an application of the new technique, we present a nearly-linear time approximation algorithm for uncapacitated minimum-cost flow on undirected graphs. Given an undirected graph with m edges labelled with costs, and n vertices labelled with demands, the algorithm takes $\epsilon^{-2}m^{1+o(1)}$-time and outputs a flow routing the demands with total cost at most $(1+\epsilon)$ times larger than minimal, along with a dual solution proving near-optimality. The generalized preconditioner is obtained by embedding the cost metric into ℓ_1, and then considering a simple hierarchical routing scheme in ℓ_1 where demands initially supported on a dense lattice are pulled from a sparser lattice by randomly rounding unaligned coordinates to their aligned neighbors. Analysis of the generalized condition number for the corresponding preconditioner follows that of the classical multigrid algorithm for lattice Laplacian systems.

1 Introduction

A fundamental problem in optimization theory is that of finding solutions x of minimum-norm $\|x\|$ to rectangular linear systems $\mathbf{A}x = \mathbf{b}$. Such solvers have extensive applications, due to the fact that many practical optimization problems reduce to minimum norm problems. When the norm is Euclidean, classical iterative solvers such as steepest descent and conjugate gradient methods produce approximately optimal solutions with residual error $\|\mathbf{A}x - \mathbf{b}\|$ exponentially small in iteration count, The rate of exponential convergence depends strongly on the *condition number* of the square matrix $\mathbf{A}\mathbf{A}^*$.

Therefore, the algorithm design process for such problems typically consists entirely of efficiently constructing *preconditioners*: easily computable left-cancellable operators $\mathbf{P}$, for which the transformed problem $\mathbf{P}\mathbf{A} = \mathbf{P}\mathbf{b}$ is well-conditioned so iterative methods converge rapidly. In a seminal work, Spielman and Teng[17] present a nearly-linear-time algorithm to construct preconditioners with condition number polylogarithmic in problem dimension for a large class of operators $\mathbf{A}$, including Laplacian systems and others arising from discretization of elliptic PDEs.

There are many fundamental optimization problems that can be expressed as minimum-norm problems with respect to more general, non-Euclidean norms, including some statistical inference problems and network flow problems on graphs. In particular, the fundamental *maximum-flow* and *uncapacitated minimum-cost flow* problems reduce respectively to ℓ_∞ and ℓ_1 minimum-norm problems in the undirected case.

For such problems, there is presently no known black-box iterative solver analagous to those for ℓ_2 that converge exponentially with rate independent of problem dimension. Therefore, one is forced to choose between interior point methods to obtain zero or exponentially small error[4], at the cost of iteration count depending strongly on problem dimension, or alternative methods that have error only polynomially small with respect to iterations[15, 9, 14].

*Research supported by NSF Grant CCF-1410022

An important idea implicitly used by classical iterative ℓ_2 solvers is *residual recursion*: reducing the error of an approximate solution x by recursively applying the solver to $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{A}x$. Such recursion is implicit in those solvers due to the fact that in ℓ_2, the map $\mathbf{b} \mapsto \mathbf{x}_{\mathsf{opt}}$ is linear, with $\mathbf{x}_{\mathsf{opt}} = (\mathbf{A}\mathbf{A}^*)^+\mathbf{b}$. Therefore, the classical methods are effectively recursing on *every* iteration, with no distinction between iterations and recursive solves. In general norms, an $\mathbf{x}_{\mathsf{opt}}$ may not be unique, and there is generally no linear operator mapping $\mathbf{b}$ to some $\mathbf{x}_{\mathsf{opt}}$.

In this paper, we make two primary contibutions. Our first contribution is to extend analysis of residual recursion to arbitrary norms. We present a general and modular framework for efficiently solving minimum-norm problems by composing simple well-known base solvers. The main tool is the *composition lemma*, describing the approximation parameters of a solver composed via residual recursion of two black-box solvers. The resulting parameters depend on a generalization of the condition number to arbitrary norms; therefore, much like ℓ_2, the algorithm design process is reduced to constructing good *generalized preconditioners*. Our second contribution is to present, as a practical application of those tools, a nearly-linear time approximation algorithm for the uncapacitated minimum-cost flow problem on undirected graphs. Having established the former framework, the latter algorithm is entirely specified by the construction of a generalized preconditioner, and analyzed by bounding its generalized condition number.

Our framework builds upon earlier work[16], where we leveraged ideas introduced by Spielman and Teng to obtain nearly-linear time approximation algorithms for undirected maximum flow. We have since realized that some of the techniques applied to extend ℓ_2-flows to ℓ_∞-flows are in no way specific to flows, and indeed the composition framework presented here is a generalization of those ideas to non-Euclidean minimum-norm problems. Moreover, a significantly simpler max-flow algorithm may be recovered using black-box solvers and general composition as presented here, requiring only the truly flow-specific part of the earlier work: the *congestion approximator*, now understood to be a specific instance of a more widely-applicable *generalized preconditioner*.

We proceed to discuss those contributions separately in more detail, and outline the corresponding sections of the paper. We also discuss some related work.

1.1 Minimum Norm Problems

In section 2, we define the minimum norm problem more precisely, and introduce a useful bicriteria notion of approximation we call (α, β)-solutions. In that notion, α quantifies how much $\|x\|$ relatively exceeds $\|x_{\mathsf{opt}}\|$, and β quantifies the residual error $\|\mathbf{A}x - \mathbf{b}\|$, with an appropriate relative scale factor. We may succinctly describe some frequently used algorithms as (α, β)-solvers; we discuss some examples in section 4, including steepest descent and conjugate gradient for ℓ_2, and *multiplicative weights* for ℓ_1 and ℓ_∞.

In section 3, we present the *composition lemma*, which shows that by composing a (α_1, β_1)-solver with a (α_2, β_2)-solver in a black-box manner, we obtain a (α_3, β_3)-solver with different parameter tradeoffs. The composed algorithm's parameters (α_3, β_3) depends crucially on a generalization of the classical condition number to rectangular matrices in general norms; Demko[7] provides essentially the exact such generalization needed, so we briefly recall that definition before stating and proving the composition lemma. The proof of the composition lemma is not difficult; it follows quite easily once the right definitions of (α, β)-solutions and a generalized condition number have been established. Nevertheless, the modularity of the composition lemma proves to quite useful. By recursively composing a solver with itself t times, we may decrease β exponentially with t, at the cost of a single *fixed* increase in α. By composing solvers with different parameters, better parameters may be obtained than any single solver alone. A particularly useful example uses a $(M, 0)$-solver, where M is uselessly-large on its own, to terminate a chain of t solvers and obtain zero error. As the final step in the chain, we get all of the benefits (zero error), and almost none of the costs, as it contributes $M2^{-t}$ instead of M to the final solver.

By composing the existing solvers discussed in section 4, we obtain solvers for ℓ_1 and ℓ_∞ problems with residual error exponentially small in t. We state the final composed algorithms parameters. As a result, the algorithm design problem for such problems is reduced to that of designing a *generalized preconditioner* for $\mathbf{A}$. We briefly discuss such preconditioning in section 5.

1.2 Minimum Cost Flow

The second part of this paper applies generalized preconditioning to solve a fundamental problem in network optimization. In the *uncapacitated minimum-cost flow problem* on undirected graphs, we are given a connected graph G with m edges, each annotated with a cost, and a specified demand at each vertex. The

problem is to find an edge flow with vertex divergences equal to the specified demands, of minimal total cost. Our main result is a randomized algorithm that, given such a problem, takes $\epsilon^{-2}m^{1+o(1)}$ time and outputs a flow meeting the demands with cost at most $(1+\epsilon)$-times that of optimal. We define various flow-related terms and give a more precise statement of the problem and our result in section 6.

We describe and analyze the algorithm in section 7. Having built up our general tools in earlier sections, our task is reduced to the design and analysis of a generalized preconditioner for min-cost flow. We begin by interpreting the edge costs as lengths inducing a metric on the graph, and then apply Bourgain's small-distortion embedding[3] into ℓ_1. That is, by paying a small distortion factor in our final condition number, we may focus entirely on designing a preconditioner for min-cost flow in ℓ_1 space. The construction of our precondtitioner is based on an extremely simple hierarchical routing scheme, inspired by a combination of the multigrid algorithm for grid Laplacians, and the Barnes-Hut algorithm for n-body simulation[2]. The routing scheme proceeds on a sequence of increasingly dense lattices $V_0, V_1, \ldots$ in ℓ_1, where V_t are the lattice points with spacing 2^{-t}, and the input demands are specified on the densest level V_T. Starting at level $t = T$, the routing scheme sequentially eliminates the demand supported on $V_t \setminus V_{t-1}$ by pulling the demand from level $t-1$ using a random-rounding based routing: each vertex $x \in V_t$ picks a random nearest neighbor in V_{t-1}, corresponding to randomly rounding its unaligned coordinates, and then pulls a fractional part of its demand from that neighbor via any shortest path. Afterwards, all demands on $V_t \setminus V_{t-1}$ are met, leaving a reduced problem with demands supported on V_{t-1}. The scheme then recurses on the reduced problem. At level 0, the demands are supported on the corners of a hypercube, for which the simple routing scheme of pulling all demand from a uniformly random corner performs well-enough. Having described the routing scheme, we need not actually carry it out. Instead, our preconditioner crudely estimates the cost of that routing. Furthermore, while the routing has been described on an infinite lattice, we'll observe that if the demands are only initially supported on a small set of n vertices, the support remains small throughout.

1.3 Related Work

Spielman and Teng present a nearly-linear time algorithm for preconditioning and solving symmetric diagonally dominant matrices, including those of graph Laplacians[17]. The ideas introduced in that work have led to breakthroughs in various network cut and flow algorithms for the case of undirected graphs. Christiano *et. al.* apply the solver as a black-box to approximately solve the maximum flow problem in $\Theta(m^{4/3})$ time. Madry[13] opens the box and uses the ideas directly to give a family of algorithms for approximating cut-problems, including a $m^{o(1)}$-approximation in $m^{1+o(1)}$ time. Both the present author[16] and Kelner *et. al.*[11] combine madry's ideas with ℓ_∞ optimization methods to obtain $(1-\epsilon)$-approxmations to maximum flow in $m^{1+o(1)}\epsilon^{-\mathcal{O}(1)}$ time. While those two algorithms use similar ideas, the actual path followed to achieve the result is rather different. Kelner *et. al.* construct an *oblivious routing scheme* that is $n^{o(1)}$-competetive, and then show how to use it to obtain a flow. The algorithm maintains a demand-respecting flow at all times, and minimizes a potential function measuring edge congestion. In contrast, our earlier algorithm short-cuts the need to explicitly construct the routing scheme by using Madry's construction directly, maintaining a flow that is neither demand nor capacity respecting, aiming to minimize a certain potential function that measures both the congestion and the demand error.

For min-cost flow, earlier work considers the more general directed, capacitated case, with integer capacities in $\{1, \ldots, U\}$. The $\Theta(nm \log\log U)$-time double-scaling algorithm of Ahuja, Goldberg, Orlin, and Tarjan[1] remained the fastest algorithm for solving min-cost flow for 25 years between its publication and the Laplacian breakthrough. Shortly after that breakthrough, Daitch and Spielman[5] showed how to use the Laplacian solver with interior-point methods to obtain an $\Theta(m^{3/2}\log^2 U)$-time algorithm. Lee and Sidford further reduce this to $\Theta(m\sqrt{n}\log^{\mathcal{O}(1)} U)$ by a general improvement in interior point methods[12]. We are not aware of prior work

2 Approximate Solutions

Let $\mathcal{X}, \mathcal{Y}$ be finite dimensional vector spaces, where $\mathcal{X}$ is also a Banach space, and let $\mathbf{A} \in \mathrm{Lin}(\mathcal{X}, \mathcal{Y})$ be fixed throughout this section. We consider the problem of finding a minimal norm pre-image of a specified $\mathbf{b}$ in the image of $\mathbf{A}$; that is, finding $x \in \mathcal{X}$ with $\mathbf{A}x = \mathbf{b}$ and $\|x\|_{\mathcal{X}}$ minimal.

Let $\mathbf{x}_{\mathsf{opt}}$ be an exact solution, with $\mathbf{A}\mathbf{x}_{\mathsf{opt}} = \mathbf{b}$, and $\|\mathbf{x}_{\mathsf{opt}}\|_{\mathcal{X}}$ minimal. There are multiple notions of

approximate solutions for this problem. The most immediate is $x \in \mathcal{X}$ with $\mathbf{A}x = \mathbf{b}$ and

$$\frac{\|x\|}{\|\mathbf{x}_{\mathsf{opt}}\|} \leq \alpha \tag{1}$$

. We call such x an $(\alpha, 0)$-solution; we shall be interested in finding $(1+\epsilon, 0)$-solutions for small ϵ. A weaker notion of approximation is obtained by further relaxing $\mathbf{A}x = \mathbf{b}$ to $\mathbf{A}x \approxeq \mathbf{b}$. Quantifying that requires more structure on $\mathcal{Y}$, so let us further assume $\mathcal{Y}$ to also be a Banach space. In that case, we say x is an (α, β)-solution if equation 1 holds and

$$\frac{\|\mathbf{A}x - \mathbf{b}\|}{\|\mathbf{A}\|\|\mathbf{x}_{\mathsf{opt}}\|} \leq \beta \tag{2}$$

The practical utility of such solutions depends on the application. However, the weaker notion has the distinct advantage of being approachable by a larger family of algorithms, such as penalty and dual methods, by avoiding equality constraints. We discuss such existing algorithms in section 4, but mention that they typically yield $(1, \epsilon)$-solutions after some number of iterations. For ℓ_2, the iteration dependency on ϵ is $\mathcal{O}(\log(1/\epsilon))$, while for some more general norms it is $\epsilon^{-\mathcal{O}(1)}$.

3 Recursive Composition and Generalized Condition Numbers

We now consider how to trade an increase in α for a decrease in β. Residual recursion suggests a natural strategy: after finding an (α, β)-solution, recurse on $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{A}x$. More precisely, we define the *composition* of two algorithms as follows.

Definition 3.1. *Let F_i be an (α_i, β_i)-algorithm for $\mathbf{A}$, for $i \in \{1, 2\}$. The* composition *$F_2 \circ F_1$ takes input $\mathbf{b}$, and first runs F_1 on $\mathbf{b}$ to obtain x. Next, setting $\tilde{\mathbf{b}} = \mathbf{b} - \mathbf{A}x$, it runs F_2 on input $\tilde{\mathbf{b}}$ to obtain $\tilde{x}$. Finally, it outputs $x + \tilde{x}$.*

Success of composition depends on $\|\tilde{\mathbf{b}}\|_{\mathcal{Y}}$ being small implying $\tilde{\mathbf{b}}$ has a small-norm pre-image. The extent to which that is true is quantified by the *condition number* of $\mathbf{A}$. The condition number in ℓ_2 may be defined several ways, resulting in the same quantity. When generalized to arbitrary norms, those definitions differ. We recall two natural definitions, following Demko[7].

Definition 3.2. *The* non-linear condition number *of $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$ is*

$$\tilde{\kappa}_{\mathcal{X}\to\mathcal{Y}}(\mathbf{A}) = \min\left\{ \frac{\|\mathbf{A}\|_{\mathcal{X}\to\mathcal{Y}}\|x\|_{\mathcal{X}}}{\|\mathbf{A}x\|_{\mathcal{Y}}} : \mathbf{A}x \neq 0 \right\}$$

The linear condition number *is defined by*

$$\kappa_{\mathcal{X}\to\mathcal{Y}}(\mathbf{A}) = \min\left\{\|\mathbf{A}\|_{\mathcal{X}\to\mathcal{Y}}\|\mathbf{G}\|_{\mathcal{Y}\to\mathcal{X}} : \mathbf{G} \in \mathrm{Lin}(\mathcal{Y}, \mathcal{X}) : \mathbf{AGA} = \mathbf{A}\right\}$$

Of course, $\tilde{\kappa} \leq \kappa$. Having defined the condition number, we may now state how composition affects the approximation parameters.

Theorem 3.3 (Composition)**.** *Let F_i be an $(\alpha_i, \beta_i/\tilde{\kappa})$-algorithm for $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$, where $\mathbf{A}$ has non-linear condition number $\tilde{\kappa}$ Then, the composition $F_2 \circ F_1$ is an $(\alpha_1 + \alpha_2\beta_1, \beta_1\beta_2/\tilde{\kappa}$-algorithm for the same problem.*

Before proving lemma 3.3, we state two useful corollaries. The first concerns the result of recursively composing a $(\alpha, \beta/\tilde{\kappa})$-algorithm with itself.

Corollary 3.4. *Let F be a $(\alpha, \beta/\tilde{\kappa})$-algorithm for $\beta < 1$. Let F^t be the sequence formed by iterated composition, with $F^1 = F$, $F^{t+1} = F^t \circ F$. Then, F^t is a $(\alpha/(1-\beta), \beta^t/\tilde{\kappa})$-algorithm.*

Proof. By induction on t. To start, an $(\alpha, \beta/\tilde{\kappa})$-algorithm is trivially a $(\alpha/(1-\beta), \beta/\tilde{\kappa})$-algorithm. Assuming the claim holds for F^t, lemma 3.3 implies F^{t+1} is a $(\alpha + \alpha\beta/(1-\beta), \beta^{t+1}/\tilde{\kappa})$-algorithm. □

We observe that to obtain $(1 + \mathcal{O}(\epsilon), \delta)$-solution, only the first solver in the chain need be very accurate.

Corollary 3.5. *Let F be a $(1+\epsilon, \epsilon/2\tilde{\kappa}$-solver and G be a $(2, 1/2\tilde{\kappa})$-solver. Then, $G^t \circ F$ is a $(1+5\epsilon, \epsilon 2^{-t-1}/\tilde{\kappa}$-solver.*

For some problems, there is a very simple $(M, 0)$-solver known. A final composition with that solver serves to elimate the error.

Corollary 3.6. *Let F be a $(1+\epsilon, \epsilon\delta/\tilde{\kappa})$-solver and G be a $(M, 0)$-solver. Then $G \circ F$ is a $(1+\epsilon(1+\delta M), 0)$-solver.*

3.1 Proof of lemma 3.3

Since F_1 is an $(\alpha_1, \beta_1/\tilde{\kappa})$-algorithm, we have

$$\begin{aligned} \|x\| &\leq \alpha_1 \|\mathbf{x}_{\mathsf{opt}}\| \\ \frac{\|\tilde{\mathbf{b}}\|}{\|\mathbf{A}\|} &\leq \frac{\beta_1}{\tilde{\kappa}} \|\mathbf{x}_{\mathsf{opt}}\| \end{aligned}$$

By the definition of $\tilde{\kappa}$,

$$\|\tilde{\mathbf{x}}_{\mathsf{opt}}\| \leq \tilde{\kappa} \frac{\|\tilde{\mathbf{b}}\|}{\|\mathbf{A}\|} \leq \beta_1 \|\mathbf{x}_{\mathsf{opt}}\|$$

Since F_2 is an $(\alpha_2, \beta_2/\tilde{\kappa})$-algorithm,

$$\begin{aligned} \|\tilde{x}\| &\leq \alpha_2 \|\tilde{\mathbf{x}}_{\mathsf{opt}}\| \leq \alpha_2 \beta_1 \|\mathbf{x}\|_{\mathsf{opt}} \\ \frac{\|\mathbf{A}\tilde{x} - \tilde{\mathbf{b}}\|}{\|\mathbf{A}\|} &\leq \frac{\beta_2}{\tilde{\kappa}} \|\tilde{\mathbf{x}}_{\mathsf{opt}}\| \leq \frac{\beta_1 \beta_2}{\tilde{\kappa}} \end{aligned}$$

The conclusion follows from $\|x + \tilde{x}\| \leq \|x\| + \|\tilde{x}\|$.

4 Solvers for ℓ_p

The recursive composition technique takes existing (α, β)-approximation algorithms and yields new algorithms with different approximation parameters. In this section, we discuss the parameters of existing well-known base solvers.

Let $\mathbf{A} : \mathbb{R}^m \to \mathbb{R}^n$ be linear and fixed throughout this section. To concisely describe results and ease comparison, for a norm $\|\cdot\|_{\mathcal{X}}$ on $\mathbb{R}^m$ and a norm $\|\cdot\|_{\mathcal{Y}}$ in $\mathbb{R}^n$, we write $(\alpha, \beta)_{\mathcal{X} \to \mathcal{Y}}$ to denote an (α, β) solution with respect to $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$.

The classical ℓ_2 solvers provide a useful goalpost for comparison.

Theorem 4.1 ([4])**.**

The steepest-descent *algorithm produces a* $(1, \delta)_{2\to 2}$*-solution after* $\mathcal{O}(\kappa_{2\to 2}(\mathbf{A})^2 \log(1/\delta))$ *simple iterations*

The conjugate gradient *algorithm produces a* $(1, \delta)_{2\to 2}$*-solution after* $\mathcal{O}(\kappa_{2\to 2}(\mathbf{A}) \log(1/\delta))$ *simple iterations.*

There are many existing algorithms for p-norm minimization, and composing them yields algorithms with exponentially small error. For this initial manuscript, we simply state the ℓ_1 version required for min-cost flow. The general cases consist of describing the many existing algorithms[14] in terms of (α, β)-solvers.

Theorem 4.2. *There is a* $(1 + \epsilon, \delta)_{1\to 1}$*-solver with simple iterations totalling*

$$\mathcal{O}\left(\tilde{\kappa}_{1\to 1}(\mathbf{A})^2 \log(m) \left(\epsilon^{-2}\right) + \log(\delta^{-1})\right)\right)$$

Theorem 4.3. *There is a* $(1 + \epsilon, \delta)_{\infty\to\infty}$*-solver with simple iterations totalling*

$$\mathcal{O}\left(\tilde{\kappa}_{\infty\to\infty}(\mathbf{A})^2 \log(n) \left(\epsilon^{-2} + \log(\delta^{-1})\right)\right)$$

The preceding theorems follow by combining the *multiplicative weights algorithm*[15, 9] with the composition lemma. The former algorithm applies in a more general online setting; for our applications, the algorithm yields the "weak" approximation algorithms that shall be composed.

A *support oracle* for a compact-convex set $\mathbf{C}$ takes input y and returns some $x \in \mathbf{C}$ maximizing $x \cdot \mathbf{C}$. Let $\triangle_n$ be the unit simplex in $\mathbb{R}^n$ (i.e., the convex hull of the standard basis).

The multiplicative weights method finds approximate solutions to the saddle point problem,

$$\max_{w \in \triangle_n} \min_{z \in \mathbf{C}} w \cdot z$$

In particular, it obtains an additive ϵ-approximation in $\mathcal{O}(\rho^2 \epsilon^{-2} \log n)$ iterations with each iteration taking $\mathcal{O}(n)$ time plus a call to a support oracle for $\mathbf{C}$.

Let $\mathbf{B}_p$ be the unit ball in ℓ_p. A point $w \in \mathbf{B}_1$ is represented as $w = w_+ - w_-$ for $(w_+, w_-) \in \triangle_{2n}$. The base solvers follow by considering the saddle-point problems,

$$\max_{y^* \in \mathbf{b}_1} \min_{x \in \mathbf{b}_\infty} \quad y^* \cdot (\mathbf{A}x - \mu \mathbf{b})$$
$$\max_{x \in \mathbf{b}_1} \min_{y^* \in \mathbf{b}_\infty} \quad y^* \cdot (\mathbf{A}x - \mu \mathbf{b})$$

where μ is scaled via binary-search as needed. The search adds a factor $\mathcal{O}(\log \tilde{\kappa})$ to the complexity. This factor may be avoided by using ℓ_p-norm for $p = \log n$ and $p = (\log m)/\epsilon$ regularization, respectively, instead.

5 Generalized Preconditioning

Let $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$ be linear. The minimum-norm problem for $\mathbf{A}$ does not intrinsically require $\mathcal{Y}$ to be normed; nor does the definition of a $(\alpha, 0)$-solution. Thus, an algorithm designer seeking to solve a class of minimum norm problems may freely choose how to norm $\mathcal{Y}$. The best choice requires balancing two factors. First, the norm should be sufficiently "simple" that $(1, \epsilon)$ solutions for polynomially-small ϵ are easy to find. Second, the norm should be chosen so that $\tilde{\kappa}(\mathbf{A})_{\mathcal{X} \to \mathcal{Y}}$ is small. If only the latter constraint existed, the ideal choice would be the *optimal* $\mathbf{A}$ *pre-image norm*,[1]

$$\|\mathbf{b}\|_{\mathsf{opt}(\mathbf{A})} = \min\{\|x\|_{\mathcal{X}} : \mathbf{A}x = \mathbf{b}\}$$

By construction, $\tilde{\kappa}(\mathbf{A})_{\mathcal{X} \to \mathsf{opt}(\mathbf{A})} = 1$. Of course, simply evaluating that norm is equivalent to the original problem we aim to solve. It follows that $\mathcal{Y}$ should be equipped with the relatively closest norm to $\mathsf{opt}(\mathbf{A})$ that can be efficiently minimized.

Following the common approach in ℓ_2, when $\mathcal{X} = \ell_p$, we propose to choose a *generalized pre-conditioner* $\mathbf{P} : \mathcal{Y} \to \ell_p$ injective and consider the norm $\|\mathbf{P}\mathbf{b}\|_p$. Again, $\mathbf{P}$ should be chosen such that $\tilde{\kappa}(\mathbf{P}\mathbf{A})_{p \to p}$ is small, and $\mathbf{P}, \mathbf{P}^*$ are easy to compute.

6 Graph Problems

In this section we discuss some applications of generalized preconditioning to network flow problems. Let $\mathcal{G} = (V, E)$ be an undirected graph with n vertices and m edges. While undirected, we assume the edges are oriented arbitrarily. We denote by $\mathcal{V}^*, \mathcal{E}^*$ the spaces of real-valued functions $f : V \to \mathbb{R}$ and $f : E \to \mathbb{R}$ on vertices and edges, respectively. We denote by $\mathcal{V}, \mathcal{E}$ the corresponding dual spaces of *demands* and *flows*. For $S \subseteq V$, we write $\mathbf{1}_S \in \mathcal{V}^*$ for the indicator function on S.

The *discrete derivative operator* $\mathcal{D}^* : \mathcal{V}^* \to \mathcal{E}^*$ is defined by $(\mathcal{D}^* f)(xy) = f(y) - f(x)$ for each oriented edge xy. The constant function has zero derivative, $\mathcal{D}^* \mathbf{1}_V = 0$. The (negative) adjoint $-\mathcal{D} : \mathcal{E} \to \mathcal{V}$ is the *discrete divergence operator*; for a flow $\jmath \in \mathcal{E}$ and vertex $x \in V$, $-(\mathcal{D}\jmath)(x)$ is the net quantity being transported away from vertex x by the flow.

A *single-commodity flow problem* is specified by a *demand vector* $\mathbf{b} \in \mathcal{V}$, and requires finding a flow $\jmath \in \mathcal{E}$ satisfying $\mathcal{D}\jmath = \mathbf{b}$, minimizing some cost function. When the cost-function is a norm on $\mathcal{E}$, the problem is a minimum-norm pre-image problem. Assuming G is connected, $\mathbf{b}$ has a pre-image iff the total demand $\mathbf{1}_V \cdot \mathbf{b}$ is zero. We hereafter assume G to be connected and total demands equal to zero.

Having established the generalized preconditioning framework, we may now design fast algorithms for fundamental network flow problems by designing generalized preconditioners for the corresponding minimum-norm problems. As a result, we obtain nearly-linear time algorithms for max-flow and uncapacitated min-cost flow in undirected graphs. The max-flow algorithm was presented previously[16], and involved a combination of problem-specific definitions (e.g. *congestion approximators*) and subroutines (e.g. gradient-based ℓ_∞ minimization). The brief summary in subsection 6.1 shows how the same result may be immediately obtained by combining theorem **??** with only small part of the earlier flow-specific work[16]: the actual construction of the preconditioner.

In section 7, we present a nearly-linear time algorithm for uncapacitated min-cost flows, another fundamental network optimization problem. Once again, all that is needed is description and analysis of the preconditioner.

[1] We remark this is only defined on the image space of $\mathbf{A}$, rather than the entire codomain. However, we only consider $\mathbf{b}$ in that image space

6.1 Max-Flow

A *capacitated graph* associates with each edge e, a *capacity* $\mathbf{c}(e) > 0$. Given capacities, we define the *capacity* norm on $\mathcal{E}^*$ by $\|f\|_{\mathbf{c}*} = \sum_e \mathbf{c}(e)|f(e)|$; the dual *congestion norm* on $\mathcal{E}$ is defined by $\|\jmath\|_{\mathbf{c}} = \max_e \frac{|\jmath(e)|}{\mathbf{c}(e)}$.

The *boundary capacity* of a set $S \subseteq \mathcal{V}$ is $\|\mathcal{D}^*\mathbf{1}_S\|_c$, which we abbreviate as $\mathbf{c}(\partial S)$.

The single-commodity flow problem with capacity norm is *undirected maximum flow problem*, and is a fundamental problem in network design and optimization. Following the work of Spielman and Teng for ℓ_2-flows[17], and Madry[13] for cut problems, the author has obtained a nearly-linear time algorithm for maximum flow[16]. We have since realized that only certain parts of the ideas introduced in that paper are actually specific to maximum-flow, and the present work has been obtained through generalizing the other parts. The (previously presented as) flow-specific composition and optimization parts of that paper are entirely subsumed by the earlier sections of this paper. The truly flow-specific part of that work required to complete the algorithm is the *generalized preconditioner* construction, which we have previously called a *congestion-approximator*. We briefly describe the main ideas of the construction, and some simple illustrative examples. For the full details, we refer the reader to the third section of [16].

Let $\mathfrak{F} \subseteq 2^{\mathcal{V}}$ be a family of subsets. Such a family, together with a capacity norm on $\mathcal{E}$, induces a *cut-congestion* semi-norm on $\mathcal{V}$ via $\|b\|_{\mathbf{c}(\mathfrak{F})} = \max_{S \in \mathfrak{F}} \frac{|\mathbf{1}_S \cdot b|}{\mathbf{c}(\partial S)}$. That is, for each set in the family, consider the ratio of the total aggregate demand in that set to the total capacity of all edges entering that set. If the indicator functions of sets in $\mathfrak{F}$ span $\mathcal{V}$, then it is a norm. Note also that we may write $\|b\|_{c(\mathfrak{F})} = \|\mathbf{P}b\|_\infty$, where $\mathbf{P} : \mathbb{R}^n \to \mathbb{R}^{\mathfrak{F}}$ is defined by $(\mathbf{P}b)_S = \frac{\mathbf{1}_S \cdot b}{\mathbf{c}(\partial S)}$.

If $\mathfrak{F}$ is taken to be the full power set $\mathfrak{F} = 2^{\mathcal{V}}$, the *max-flow, min-cut theorem* is equivalent to the statement that the non-linear condition number $\tilde{\kappa}_{\mathbf{c} \to \mathbf{c}(\mathfrak{F})}(\mathcal{D}^*)$ is exactly one. If $\mathfrak{F}$ is taken to be the family of singleton sets, the non-linear condition number is the inverse of the *combinatorial conductance* of $\mathcal{G}$. That is, for high-conductance graphs, composition yields a fast algorithm requiring only a simple *diagonal* preconditioner. This is closely analagous to the situation for ℓ_2, where diagonal preconditioning yields fast algorithms for graphs of large *algebraic conductance*.

A particularly illustrative example consists of a unit-capacity $2^t \times 2^t$ 2-D grid. With low ($\Theta(2^{-t}) = \Theta(1/\sqrt{n})$ conductance, the singleton family $\mathfrak{F}$ performs poorly. A simple but dramatically better family consists of taking $\mathfrak{F}$ to be all power-of-two size, power-of-two aligned subgrids. The latter family yields *linear* condition number $\mathcal{O}(t) = \mathcal{O}(\log n)$. Furthermore, computing the aggregate demand in all such sets requires only $\mathcal{O}(n)$ time, with each sub-grid aggregate equal to the sum of its four child aggregates. This algorithm is analagous to an ℓ_∞ version of *multi-grid*[**?**]. Note that the indicator "step" functions $\mathbf{1}_S$ are different from the bilinear "pyramid" functions used by multigrid.

Theorem 6.1 ([16])**.** *There is an algorithm that given a capacitated graph $(G, \mathbf{c})$ with n vertices and m edges, takes $m^{1+o(1)}$ time and outputs a data structure of size $n^{1+o(1)}$ that efficiently represents a preconditioner $\mathbf{P}$ with $\kappa(\mathbf{P}\mathcal{D})_{\mathbf{c}\to\infty} \leq n^{o(1)}$. Given the data structure, $\mathbf{P}$ and $\mathbf{P}^*$ can be applied in $n^{1+o(1)}$ time.*

Let $\mathbf{C} : \mathbb{R}^m \to \mathcal{E}$ be the diagonal capacity matrix; $(\mathbf{C}x)(e) = \mathbf{c}(e)x(e)$. The nearly-linear time algorithm follows by using the preconditioner $\mathbf{P}$ in the preceding theorem, applying theorem **??** to the problem of minimizing $\|x\|_\infty$ subject to $\mathbf{P}\mathcal{D}\mathbf{C}x = \mathbf{b}$. A final minor flow-specific part is needed to achieve zero error: the simple $(m, 0)$-approximation algorithm that consists of routing all flow through a maximum-capacity spanning tree.

7 Uncapacitated Min-Cost Flow

In this section, we present a nearly-linear time algorithm for the uncapacitated minimum-cost flow problem on undirected graphs. A *cost graph* associates a *cost* $\mathbf{l}(e) > 0$ with each edge e in G. The *cost* norm on the space of flows $\mathcal{E}$ is defined by $\|\jmath\|_{\mathbf{l}} = \sum_e \|\jmath(e)\|\mathbf{l}(e)$. The *undirected uncapacitated minimum-cost flow problem* is specified by the graph G, lengths ℓ, and demands $\mathbf{b} \in \mathcal{V}$, and consists of finding $\jmath \in \mathcal{E}$ with $\mathcal{D}\jmath = \mathbf{b}$ minimizing $\|\jmath\|_{\mathbf{l}}$. Note that, unlike max-flow, the *single-source* uncapacitated min-cost flow problem is significantly easier than the distributed source case: the optimal solution is to route along a single-source shortest path away from the source. The problem only becomes interesting with distributed sources and sinks.

As a new application of generalized preconditioning, we obtain a nearly-linear time algorithm for approximately solving this problem.

Theorem 7.1. *There is a randomized algorithm that, given a length-graph $G = (V, E, \mathbf{l})$, takes $\frac{m^{1+o(1)}}{\epsilon^2}$ time and outputs a $(1+\epsilon, 0)$-solution to the undirected uncapacitated min-cost flow problem.*

As expected, the bulk of the algorithm lies in constructing a good preconditioner. After introducing some useful definitions and lemmas, we spend the majority of this section constructing and analyzing a preconditioner for this problem.

Theorem 7.2. *There is a randomized algorithm that, given a length-graph G, takes $\mathcal{O}(m \log^2 n + n^{1+o(1)})$ time and outputs a $n^{1+o(1)} \times n$ matrix $\mathbf{P}$ with $\kappa_{\mathbf{l}to\infty}(\mathbf{P}\mathcal{D}) \leq n^{o(1)}$. Every column of $\mathbf{P}$ has $n^{o(1)}$ non-zero entries.*

Dual to cost is the *stretch* norm on $\mathcal{E}^*$ is $\|f\|_{\mathbf{l}*} = \max_e |f(e)|/\mathbf{l}(e)$. We say a function $\phi \in \mathcal{V}$ is *L-Lipschitz* iff $\|\mathcal{D}^*\phi\|_{\mathbf{l}*} \leq L$; that is, $|\phi(x) - \phi(y)| \leq L\mathbf{l}(xy)$ for all edges xy. The dual problem is to maximize $\phi \cdot \mathbf{b}$ over 1-Lipschitz $\phi \in \mathcal{V}$.

A length function $\mathbf{l}$ induces an intrinsic metric $\mathbf{d} : V \times V \to \mathbb{R}$, with distances determined by shortest paths. The preceding definition of Lipschitz is equivalent to $|\phi(x) - \phi(y)| \leq L\mathbf{d}(x, y)$ for all $x, y \in V$. That is, the Lipschitz constant of ϕ depends only on the metric, and is otherwise independent of the particular graph or edge lengths inducing that metric. As the dual problem depends only on the induced metric, so must the *value* of the min-cost flow. Therefore, we may unambiguously write

$$\|b\|_{\mathsf{opt}(\mathbf{d})} = \max\{\phi \cdot b : \phi \text{ is 1-Lipschitz w.r.t. } \mathbf{d}\}$$

The monotonicity of cost with distance follows immediately.

Lemma 7.3. *If $\mathbf{d}(x, y) \leq \tilde{\mathbf{d}}(x, y)$ for all $x, y \in \mathcal{V}$, then for all $b \in \mathcal{V}_{\perp\mathbf{1}}$,*

$$\|b\|_{\mathsf{opt}(\mathbf{d})} \leq \|b\|_{\mathsf{opt}(\tilde{\mathbf{d}})}$$

We recall that in some cases, the identity operator is a good preconditioner; for max-flow, such is the case in constant-degree expanders. For min-cost flow, the analagous case is a metric where distances between any pair of points differ relatively by small factors.

Definition 7.4. *Let $U \subseteq V$. We say U is* r-separated *if $\mathbf{d}(x, y) \geq r$ for all $x, y \in U$ with $x \neq y$. We say U is* R-bounded *if $\mathbf{d}(x, y) \leq R$ for all $x, y \in U$.*

Lemma 7.5. *Let $U \subset V$ be R-bounded and r-separated with respect to $\mathbf{d}$. Let $b \in \mathcal{V}_{\perp\mathbf{1}}$ be demands supported on U. Then,*

$$\frac{r}{2}\|b\|_1 \leq \|b\|_{\mathsf{opt}(\mathbf{d})} \leq \frac{R}{2}\|b\|_1$$

Proof. Note $\frac{r}{2}\|b\|_1$ is exactly the min-cost of routing b in the metric where all distances are exactly r (consider the star graph with edge lengths $r/2$, with a leaf for each $x \in U$). The inequalities follow by monotonicity. □

To complete the proof of theorem 7.1, assuming theorem 7.2, we also need a simple approximation algorithm that yields a poor approximation, but with zero error. This is easily achieved by the algorithm that routes all flow through a minimum cost (i.e., total length) spanning tree T.

Lemma 7.6. *Let T be a minimum cost spanning tree with respect to $\mathbf{l}$ on $\mathcal{G}$. Consider the algorithm that, given $b \in \mathcal{V}_{\perp\mathbf{1}}$, outputs the flow $\jmath \in \mathcal{E}$ routing b using only the edges in T. Then,*

$$\|\jmath\| \leq n\|b\|_{\mathsf{opt}(\mathbf{d})}$$

Proof. The flow on T is unique, and therefore, optimal for the metric $\tilde{\mathbf{d}}$ induced by paths restricted to lie in T. As T is a minimum cost spanning tree, $\tilde{\mathbf{d}}(x, y) \leq n\mathbf{d}(x, y)$. □

We now prove theorem 7.1. Let $\mathcal{L} : \mathbb{R}^m \to \mathcal{E}$ be the diagonal length matrix $(\mathcal{L}x)(e) = \ell(e)x(e)$, and let $\mathbf{P}$ be the preconditioner from theorem 7.2. Computing $\mathbf{P}$ takes $\mathcal{O}(m \log^2 n) + n^{1+o(1)}$ time. The minimum-cost flow problem is equivalent to the problem of minimizing $\|x\|_1$ subject to $\mathbf{P}\mathcal{D}\mathcal{L}^{-1}x = \mathbf{P}\mathbf{b}$. As $\mathbf{P}$ has $n^{1+o(1)}$ non-zero entries and $\kappa_{1\to1}(\mathbf{P}\mathcal{D}\mathcal{L}^{-1}) \leq n^{o(1)}$, the total running time is $\frac{m^{1+o(1)}}{\epsilon^2}$.

8 Preconditioning Min-Cost Flow

A basic concept in metric spaces is that of an *embedding*: $\psi : V \to \tilde{V}$, with respective metrics $\mathbf{d}, \tilde{\mathbf{d}}$. An embedding has *distortion* L if for some $\mu > 0$,

$$1 \leq \mu \frac{\tilde{\mathbf{d}}(\psi(x), \psi(y))}{\mathbf{d}(x, y)} \leq L$$

Metric embeddings are quite useful for preconditioning; if $\mathbf{d}$ embeds into $\tilde{\mathbf{d}}$ with distortion L, then the monotonicity lemma implies the relative costs of flow problems differ by a factor L. It follows that if we can construct a preconditioner $\mathbf{P}$ for the min-cost flow problem on $\tilde{\mathbf{d}}$, the same $\mathbf{P}$ may be used to precondition min-cost flow on $\mathbf{d}$, with condition number at most L-factor worse.

Our preconditioner uses embeddings in two ways. The preconditioner itself is based on a certain hierarchical routing scheme. Hierarchical routing schemes based on embeddings into tree-metrics have been studied extensively(see e.g. [8]), and immediately yield preconditioners with polylogarithmic condition number. However, we are unable to efficiently implement those schemes in nearly-linear time. Instead, we apply Bourgain's $\mathcal{O}(\log n)$-distortion embedding[3] into ℓ_1. That is, by paying an $\mathcal{O}(\log n)$ factor in condition number, we may completely restrict our attention to approximating min-cost flow in ℓ_1 space.

Theorem 8.1 (Bourgain's Embedding[3])**.** *Any n-point metric space embeds into ℓ_p space with distortion $\mathcal{O}(\log n)$. If $\mathbf{d}$ is the intrinsic metric of a length-graph $G = (V, E)$ with m edges, there is a randomized algorithm that takes $\mathcal{O}(m \log^2 n)$ time and w.h.p. outputs such an embedding, with dimension $\mathcal{O}(\log^2 n)$.*

The complexity of constructing and evaluating our preconditioner is *exponential* in the embedding's dimension. Therefore, we reduce the dimension to $\Theta(\sqrt{\log n})$, incuring another distortion factor of $\exp(\mathcal{O}(\sqrt{\log n}))$.

Theorem 8.2 (Johnson-Lindenstrauss Lemma[10, 6])**.** *Any n points in Euclidean space embed into k-dimensional Euclidean space with distortion $n^{\mathcal{O}(1/k)}$, for $k \leq \Omega(\log n)$. In particular, $k = \mathcal{O}(\log n)$ yields $\mathcal{O}(1)$ distortion, and $k = \mathcal{O}(\sqrt{\log n})$ yields $\exp(\mathcal{O}(\sqrt{\log n}))$ distortion.*

If the original points are in $d > k$ dimensions, there is a randomized algorithm that w.h.p outputs such an embedding in $\mathcal{O}(nd)$ time.

Our reduction consists of first embedding the vertices into ℓ_2 via theorem 8.1, reducing the dimension to $k = \mathcal{O}(\sqrt{\log n})$ using lemma 8.2, and then using the identity embedding of ℓ_2 into ℓ_1, for a total distortion of $\exp(O(\sqrt{\log n}))$.

To estimate routing costs in ℓ_1, we propose an extremely simple hierarchical routing algorithm in ℓ_1, and show it achieves a certain competetive ratio. The routing scheme applies to the exponentially large graph corresponding to a lattice discretization of the k-dimensional cube in ℓ_1; after describing such a scheme, we show that if the demands in that graph are supported on n vertices, the distributed routing algorithm finds no demand to route in most of the graph, and the same routing algorithm may be carried out in a $k-d$-tree instead of the entire lattice.

On the lattice, the routing algorithm consists of sequentially reducing the support set of the demands from a lattice with spacing 2^{-t-1} to a lattice with spacing 2^{-t}, until all demand is supported on the corners of a cube containing the original demand points. On each level, the demand on a vertex not appearing in the coarser level is eliminated by pushing its demand to a distribution over the aligned points nearby.

8.1 Lattice Algorithm

For $t \in \mathbb{Z}$, let $V_t = (2^{-t}\mathbb{Z})^k$ be the lattice with spacing 2^{-t}. We design a routing scheme and preconditioner for min-cost flow in ℓ_1, with initial input demands b_T supported on a bounded subset of V_T. For $t = T, T-1, \ldots$, the routing scheme recursively reduces a min-cost flow problem with demands supported on V_t to a min-cost flow problem with demands supported on the sparser lattice V_{t-1}. Given demands b_t supported on V_t, the scheme consists of each vertex $x \in V_t$ pulling its demand $b_t(x)$ uniformly from the closest points to x in V_{t-1} along any shortest path, the lengths of which are at most $k2^{-t}$. For example, a point $x = (x_1, \ldots, x_k) \in V_1$ with $j \leq k$ non-integer coordinates and demand $b_1(x)$ pulls $b_1(x)2^{-j}$ units of flow from each of the 2^j points in V_0 corresponding to rounding those j coordinates in any way. By construction, all points in $V_t \setminus V_{t-1}$ have their demands met exactly, and we are left with a reduced problem residual demands b_{t-1} supported on V_{t-1}. As we show, eventually the demands are supported entirely on the 2^k corners of a hypercube forming a

bounding box of the original support set, at which point the the simple scheme of distributing the remaining demand uniformly to all corners is a k-factor approximation.

We do not actually carry out the routing; rather, we state a crude upper-bound on its cost, and then show that bound itself exceeds the true min-cost by a factor of some κ. For the preconditioner, we shall only require the sequence of reduced demands $b_t, b_{t-1}, \ldots$, and not the flow actually routing them.

Theorem 8.3.

$$\|b_t\|_{\mathsf{opt}} \leq \sum_{s=0}^{t} k2^{-s}\|b_s\|_1 \leq 2k(t+1)\|b_t\|_{\mathsf{opt}}$$

Assuming theorem 8.3 holds, we take our preconditioner to be a matrix $\mathbf{P}$ such that,

$$\|\mathbf{P}b_T\|_1 = \sum_{t=0}^{T} k2^{-t}\|b_t\|_1$$

The proof of theorem 8.3 uses two essential properties of the routing scheme described. The first is that the cost incurred in the reduction from b_t to b_{t-1} is not much larger than the min-cost of routing b_t. The second is the reduction does not increase costs. That is, the true min-cost for routing the reduced problem b_{t-1} is at most that of routing the original demands b_t.

Lemma 8.4. *Let b_{t-1} be the reduced demands produced when the routing scheme is given b_t. Then,*

1. *The cost incurred by that level of the scheme is at most $k2^{-t}\|b_t\|_1 \leq 2k\|b_t\|_{\mathsf{opt}}$*
2. *The reduction is non-increasing with respect to min-cost: $\|b_{t-1}\|_{\mathsf{opt}} \leq \|b_t\|_{\mathsf{opt}}$*

Assuming lemma 8.4, we now prove 8.3. We argue the sum is an upper bound on the total cost incurred by the routing scheme, when applied to b_t. The left inequality follows immediately, as the true min-cost is no larger.

The total cost of the scheme is at most the cost incurred on each level, plus the cost of the final routing of b_0. Lemma 8.4(a) accounts for all terms above $s = 0$. For the final routing, the demands b_0 are supported in $\{0,1\}^k$, so the cost of routing all demand to a random corner is at most $\frac{1}{2}k\|b_0\|_1$.

For the right inequality, we observe

$$\sum_{s=0}^{t} k2^{-t}\|b_t\|_1 \leq \sum_{s=0}^{t} 2k\|b_s\|_{\mathsf{opt}}$$

The inequality follows termwise, using lemma 8.4(a) for $s \geq 1$; the $s = 0$ case follows from lemma 7.5 because V_0 is 1-separated.

8.2 Proof of Lemma 8.4

For the first part, each vertex $x \in V_t$ routes its demand $b_t(x)$ to the closest points in V_{t-1} any shortest path. As such paths have length at most $k2^{-t}$, each point x contributes at most $k2^{-t}|b_t(x)|$.

For the second part, by scale-invariance it suffices to consider $t = 1$. Moreover, it suffices to consider the case where b_1 consists of a unit demand from between two points $x, y \in V_1$ with $\|x - y\|_1 = \frac{1}{2}$. To see this, we observe any flow routing arbitrary demands b_1 consists of a sum of such single-edge flows. As the reduction is linear, if the cost of each single-edge demand-pair is not increased, then the cost their sum is not increased. By symmetry, it suffices to consider x, y with $x = (0, z)$ and $y = (\frac{1}{2}, z)$. Then, the reduction distributes x's demand to a distribution over $(0, z')$ where $z' \sim Z'$; y's demand is split over $(0, z')$ and $(1, z')$ where $z' \sim Z'$. Therefore, half of the demand at $(0, z')$ is cancelled, leaving the residual problem of routing $1/2$ unit from $(0, z')$ to $(1, z')$. That problem has cost $1/2$, by routing each fraction directly from $(0, z')$ to $(1, z')$.

References

[1] Ravindra K Ahuja, Andrew V Goldberg, James B Orlin, and Robert E Tarjan. Finding minimum-cost flows by double scaling. *Mathematical programming*, 53(1-3):243–266, 1992.

[2] Josh Barnes and Piet Hut. A hierarchical o (n log n) force-calculation algorithm. *nature*, 324(6096):446–449, 1986.

[3] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, 1985.

[4] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[5] Samuel I Daitch and Daniel A Spielman. Faster approximate lossy generalized flow via interior point algorithms. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 451–460. ACM, 2008.

[6] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of johnson and lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

[7] Stephen Demko. Condition numbers of rectangular systems and bounds for generalized inverses. *Linear Algebra and its Applications*, 78:199–206, 1986.

[8] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 448–455. ACM, 2003.

[9] Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1):79–103, 1999.

[10] William B Johnson and Joram Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemporary mathematics*, 26(189-206):1, 1984.

[11] Jonathan A Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 217–226. Society for Industrial and Applied Mathematics, 2014.

[12] Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in o (vrank) iterations and faster algorithms for maximum flow. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 424–433. IEEE, 2014.

[13] Aleksander Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 245–254. IEEE, 2010.

[14] Yu Nesterov. Smooth minimization of non-smooth functions. *Math. Program.*, 103(1):127–152, May 2005.

[15] Serge A. Plotkin, David B. Shmoys, and Eva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20:257–301, 1995.

[16] Jonah Sherman. Nearly maximum flows in nearly linear time. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 263–269. IEEE, 2013.

[17] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *CoRR*, abs/cs/0607105, 2006.